Physician social networks and treatment variations in coronary inpatient care[1]

Gary M. Fournier
Kislaya Prasad
Mary A. Burke[2]

Florida State University

May 6, 2002

[1]The authors wish to thank Peyton Young, to whom we credit the original idea for this paper.

[2]Paper prepared for the 13th Annual Conference on Health Economics, Carnegie Mellon University, June 2-3, 2002. Please address correspondence to: Gary M. Fournier, Department of Economics, Florida State University, Tallahassee, FL 32306-2180, U.S.A. Tel.(850)644-7080 Email: gfournie@coss.fsu.edu.

1 Introduction

A striking feature of data on rates of use of medical procedures is the extent of regional variation. Studies conducted at different times, in different countries, with very different insurance systems, all find patterns of variation that demonstrate regional similarities simultaneous with considerable global diversity. For example, in a comparison of hospitalization referral regions across the country, rates of coronary artery bypass grafting among Medicare enrollees varied by a factor of more than 3.5, while the rates of coronary angioplasty ranged from 2.5 to 16.9 per 1,000 enrollees (The Dartmouth Atlas of Health Care 1999). Such statistics are taken to reflect large variations in the *quality* of care, which is an expressed cause for concern. While there has been much previous examination of such findings, regional practice variations remain something of a puzzle (see Phelps, 2000). In this study, we approach this question with detailed data on coronary patients, together with an evolutionary model of choice emphasizing local social interactions.

Do the treatment variations reflect simple physician ignorance? We find this explanation unlikely given the broad dissemination of practice guidelines and quality indicators (e.g. the Cooperative Cardiovascular Project, see O'Connor et. al. 1999, and Gunnar et. al. 1990). To look for answers, we consider the complex process by which physicians acquire, interpret, and, ultimately, incorporate information into specific treatment choices. This process is inherently social, and the patterns of choice that emerge from it will reflect a particular set of social interactions. Understanding the structure and character of such interactions is therefore essential for understanding regional practice variations, and so for making any attempts to 'improve' (or render more uniform) physicians' choices.

To explain the regional variations and understand the role of social networks in their propagation, we undertake complementary empirical and theoretical inquiries. The empirical data for this study comprise a panel of physicians and their patient discharge records in Florida hospitals in 1997-1998.

We focus on coronary care patients (HCFA, 2002), and on the choice of treatment of acute myocardial infarction (AMI) and coronary atherosclerosis in particular. The data permit construction of severity of illness measures, including detailed indicators of patient-specific comorbidity conditions, (Elixhauser et al. 1998, the Charlson Comorbidity Index, Romano et. al. 1993), as well as patient age, sex, and race. Managed care has been shown to affect treatment choice (Cutler and McClellan, 1996), so we control for the diverse financial incentives for patients and physicians by separate estimation of HMO and Medicare patients in our study.

By controlling for a far broader range of factors than has been done in most previous empirical studies of the regional variation phenomenon we hope to assess the robustness of the findings. At the same time, our interest is in assembling measures of physician interaction to develop an understanding of the social processes at work in the choice of advanced treatement for heart patients.

The process by which an appropriate procedure is chosen for a patient is a complicated one. We may identify at least three stages. The process begins with a diagnosis, which itself presents several alternatives. The physician and patient may decide upon, or against, an invasive diagnostic procedure. The next stage is the treatment choice, which is likely to be made in consultation with a specialist. In the context of cardiac care, this specialist would be a cardiologist, and a decision may be whether the patient is to undergo coronary artery bypass grafting (CABG) or coronary angioplasty. Finally, the procedure is executed by a surgeon. It is possible to develop distinct explanations of regional variation by focusing on the different stages of this process.

Our own account begins from the observation that decision–makers at the treatment choice stage should rationally take into account the skill with which a procedure is likely to be executed in the final stage. In a world with spillovers and learning, the relative frequency with which procedures are used locally could affect their likelihood of success. While we highlight increasing returns in the technology, regional variation patterns can also arise in our model if physicians or patients display a desire to conform with locally prevalent choices. In either case, increasing returns could explain the emergence of uniformity in procedure choice, or resistance to new procedures. However, it does not adequately explain the pattern of global diversity. Our model adds two additional features. First, there are regional differences in the patient population, and different procedures are appropriate for different patient types. Second, we assume a local interaction structure whereby each doctor interacts with only a small number of neighbors. There is overlap in neighborhoods, and no part of the population is totally isolated. The interplay of these three features leads to the emergence of regions with uniformity of treatment choice.

Somewhat different explanations of the phenomenon focus on learning and uncertainty at the decision stage of the process (e.g. the model described by Phelps, 2000). If the relationship between the disease and treatments is not completely understood, physicians must continually learn about the correct treatment. The experience of colleagues is one possible source of information, likely to be used since it is relatively low cost. Phelps shows how regional patterns, once established, would persist if physicians use Bayes' rule to update their beliefs. It should be possible to extend our basic approach to models of information acquisition and learning about the relative efficacy of procedures.

The econometric model focuses on the various aspects of our account of the regional variation phenomenon. To determine if social networks play a significant role, we measure the influence on a given physician's choices of the rates of usage of various procedures by other physicians working at the same hospital or set of hospitals. To test for the presence of scale economies we include raw frequency counts for the procedures of interest at the hospital at which the given procedure was performed. By including these variables in our empirical analyses, together with variables representing eleven geographic districts and several patient characteristics, we gain a more complex picture of the shape of the regional variations. Notably, we find some evidence that what have been perceived as variations by region may be better explained by interactions between physicians.

In this paper, we draw upon several different lines of research. Our theoretical model follows most directly from Young and Burke, 2001, and also builds upon ideas in the evolutionary game theory literature (in particular, Young 1993, Young 1996, Kandori, Mailath and Rob 1993, Ellison 1993, Ellison and Fudenberg 1993, and Morris 2000). In the health economics literature, the phenomenon is discussed in the survey by Phelps (2000), who also describes some of the models in considerable detail. The medical literature is extensive, even after we limit attention to cardiac care. In addition to THE DARTMOUTH ATLAS (*op cit.*), the phenomenon is described in the THE DARTMOUTH ATLAS OF CARDIOVASCULAR HEALTH CARE, 1999, and in O'Connor et. al. (1999), Pilote et. al. (1995), and numerous papers by Wennberg and coauthors (e.g. see the citations in Phelps).

The rest of the paper is organized as follows. In section 2, we describe our theoretical model and results, together with simulations of the model. Section 3 contains the empirical analysis. We begin with a discussion of the data, describe the econometric model, present results from the analysis of angiography, and of surgical interventions. Concluding remarks are in section 4.

2 A Model of Procedure Choice

The model has two essential features. Social (non-market) influences are important, and interaction among agents is local. We imagine a population of physicians, at fixed locations, interacting with a set of neighbors. The procedures used by neighbors will influence the choice of a specific physician. This may be because of a preference for conformity, but could equally well arise from spillovers in knowledge and experience. If the local pool of experience and skill with the use of a particular procedure is greater physicians are more likely to lean in favor of using it. There is a substantial body of research, starting from the work of Schelling (1971) on neighborhood segregation, which shows how the presence of interdependent preferences, spillovers, or increasing returns, can lead to much greater uniformity than is warranted by fundamentals. The empirical features of data on practice variation suggest uniformity within regions, but considerable diversity across regions. To capture this, and motivated by the model of Young and Burke (2001), we assume a local interaction structure. As we will see later, it is the combination of these two ingredients of our model which generates the regularities in the data. We assume there are population differences across regions, and find that small differences in the population profile give rise to large differences in procedure use.

The pattern of use of procedures arises as a steady state of a stochastic dynamic process. We begin by describing the dynamical system. Physicians are indexed by $\mathbb{Z}$, the set of integers. Let $\mathcal{N} \equiv \{-1, +1\}$. For each $x \in \mathbb{Z}$, $x + \mathcal{N}$ denotes the set of neighbors of x. There are two types of patients, denoted α and β, and two procedures A and B. Let $\pi_z(t, L, R)$ denote the payoff of a physician from using procedure $z \in \{A, B\}$ on a patient of type t when her neighbors use $\{L, R\}$ (L and R belong to $\{A, B\}$). For instance, $\pi_A(\alpha, \cdot, \cdot)$ could be the likelihood of success of procedure A on an α-patient, but we allow for other possibilities. We assume π_A and π_B are the same for all physicians. The essential feature of payoffs is that:

(a) Procedure A is optimal for α–patients if even one neighbor uses A.

(b) Procedure B is optimal for β–patients if even one neighbor uses B.

Three properties of payoffs can generate this feature: (1) payoffs from using a procedure increase with the number of neighbors who use the same procedure, (2) neither procedure dominates the other, and (3) for any fixed neighborhood, A yields higher payoffs when used on an α type than when used on a β type (and B yields higher payoffs when used on a β type than on an α type). However, it is not true that procedure A is always better than procedure B for an α type, nor that B is better than A for β types. For example,

$$\begin{aligned} \pi_A(\alpha, B, B) &= 0.3 \quad & \pi_A(\beta, B, B) &= 0.2 \\ \pi_A(\alpha, A, B) &= 0.4 \quad & \pi_A(\beta, A, B) &= 0.3 \\ \pi_A(\alpha, A, A) &= 0.5 \quad & \pi_A(\beta, A, A) &= 0.4 \end{aligned}$$

Similarly, for B,

$$\begin{aligned} \pi_B(\alpha, B, B) &= 0.4 \quad & \pi_B(\beta, B, B) &= 0.5 \\ \pi_B(\alpha, A, B) &= 0.3 \quad & \pi_B(\beta, A, B) &= 0.4 \\ \pi_B(\alpha, A, A) &= 0.2 \quad & \pi_B(\beta, A, A) &= 0.3 \end{aligned}$$

Patients arrive randomly at each location. Inter-arrival times are exponential with parameter λ. We will assume $\lambda = 1$.

The concentration of patient types varies by region. We partition $\mathbb{Z}$ into two regions, East and West. The negative integers constitute the West, while the non-negative integers constitute the East. The probability that a patient who arrives at any given location in the East (West) is of type α will be given by p_E (p_W). The *state* of the system is an assignment of zeros and ones to all of the integers ($\omega : \mathbb{Z} \to \{0, 1\}$). A '1' at any location indicates that the physician there used procedure A on her most recent patient. A '0' denotes the use of procedure B. We let Ω denote the set of states.

Consider a specific location $x \in \mathbb{Z}$. When a patient arrives at x, the physician makes a choice between A and B. The choice depends on the type of patient, as well as the choices made (in the recent past) by neighboring physicians. As mentioned earlier, inter-arrival times are exponential, with parameter $\lambda = 1$. We can imagine an infinite binary sequence, with the values at each location indicating the most recent choice made by the physician there. At random dates there is a transition: the value at one location changes from zero to one or vice versa. The process is a continuous time Markov chain, X_t, and we are interested in the invariant (equivalently stationary, or equilibrium) measures of this process.

Let $0 \in \Omega$ denote the state ω with $\omega(i) = 0$ for all $i \in \mathbb{Z}$. Similarly, $1 \in \Omega$ denotes the state ω with $\omega(i) = 1$ for all $i \in \mathbb{Z}$. Clearly, δ_0 and δ_1 are invariant measures. If we somehow reach the configuration 0 (or 1), the process can never escape from this state. Following Liggett (1999), we say that the process *coexists* if there is an invariant measure that is not a mixture of δ_0 and δ_1. We show that the process X_t defined above coexists by identifying such an invariant measure. That is, we identify an invariant distribution in which both procedures are used with strictly positive probability.

Define the set of states $Z \subset \Omega$ as follows: $\omega \in Z$ if there exists $m \in \mathbb{Z}$ such that $\omega(i) = 1$ for all $i < m$ and $\omega(i) = 0$ for all $i \geq m$. Z is a communicating class of the process, that is every state in Z is reached with positive probability from any other state in Z, and it is closed (once in Z, we can never escape). It is recurrent (we return to every state) but not periodic. The process restricted to Z is irreducible.

In proving the existence of an invariant distribution which has Z as its support, for simplicity we characterize such a distribution in terms of the location of the boundary point between the region in which procedure A is used and the region in which procedure B is the norm. In the proposition below, $\rho(\cdot)$ specifies the probability distribution of this boundary point. The proof is in the appendix.

Proposition 1. *Suppose $p_W > 1/2$ and $p_E < 1/2$. Then the process coexists. Specifically, there is an invariant measure ρ, with support $\mathbb{Z}$, such that*

$$\rho(m) = \frac{1}{K}\left(\frac{1-p_W}{p_W}\right)^{-m} \qquad \text{if } m < 0$$

$$\rho(m) = \frac{1}{K}\left(\frac{p_E}{1-p_E}\right)^{m} \qquad \text{if } m \geq 0.$$

K is a real number constant which can be chosen to ensure that ρ is a probability.

The proposition above tells us that the location of the boundary between procedure regions behaves according to an invariant distribution. Imagine the process as follows: each state consists of an infinite string of ones followed by infinitely many zeros, but the boundary between the two regions keeps moving around, according to the probabilities indicated by $\rho(\cdot)$. The process behaves rather like an asymmetric random walk, in which the respective probabilities in the East and in the West that a given location is the boundary differ.

In case $p_W < 1/2$ and $p_E > 1/2$, we get a similar result, only the support now consists of a string of zeros followed by ones. In case $p_W < 1/2$ and $p_E < 1/2$, the invariant distribution is δ_0. If $p_W > 1/2$ and $p_E > 1/2$, it is δ_1.

Thus the process has several invariant distributions, and we would like to identify the distribution which is most likely to be selected in the long run from randomly chosen initial conditions. To gain insight into such selection, we turn to a computational analysis. We are forced to limit attention to a finite number of locations, and consider N locations on a circle. We delimit the integers and make 1 and N neighbors. Since $\rho(\cdot)$ has support $\mathbb{Z}$, we expect that for small N one of 0 or 1 will eventually be reached, so that δ_0 or δ_1 is selected. This is indeed what we find. But for even moderately sized N we obtain the invariant distribution described in Proposition 1. For $N = 100$, we have never seen anything else in simulations. Further evidence in support of $\rho(\cdot)$ is provided by a stability analysis. When we introduce a small amount of noise in decision making (physicians occasionally experiment with procedures), δ_0 and δ_1 seem to become unstable, and we are much more likely to see ρ, even with small N. All in all, we see strong evidence to support the selection of $\rho(\cdot)$.

The results of our simulations are conveyed through graphs. We start with a random assignment of procedures to each of the N locations, and generate a decision date (arrival time of patient) for each physician. Interarrival times at each location are exponentially distributed. The probability that a patient is of a particular type differs across regions. We identify the physician with the earliest decision date, generate a patient for her, and update her procedure choice in the manner specified above. Finally, we generate the next decision date for this physician, and identify the next physician who must choose a procedure. We can then track the evolution of the procedure use at each location.

In Figure 1 we plot the frequency of use of procedure A at each location. The x–axis has $N = 100$ locations. We run the simulation for 5000 decision opportunities, and compute the proportion of times that procedure A was the recorded procedure at each location. Note that the number of actual decision opportunities at each location will be a much smaller number. The probability of arrival of α-patients is $p_W = 0.7$ in the West (locations 1–50), and $p_E = 0.3$ in the East (locations 51–100). There is no noise in this simulation. Examining Figure 1, we see that there is substantial regional variation. In the West, there is extensive use of procedure A, whereas in the East, A is infrequently used.

In Figure 2, we depict the aggregate use of procedure A at each instant that a decision opportunity arises. The decision opportunity number tracks evolution with time (only, the length of time between any two decision dates is randomly distributed). We observe that this number changes, but stays within a fairly narrow range around 50. A much better idea of the evolution of the system is obtained from Figure 3. Here we fill in a grid with two colors, gray and black (for those who cannot afford the colored luxury edition of this paper). At each decision opportunity, gray at location m indicates that the most recent procedure used there was A. At the beginning, we see a completely random pattern of procedure choices. Then, with time, we see the evolution of greater uniformity. Locations 1–50, in the West, turn completely gray, indicating adoption of procedure A. Locations 51–100 turn black, as B is adopted widely in the East.

The simulations reported in Figures 4–6 take a total of $N = 120$ locations around a circle. This is split into four regions, with thirty physicians in each location. In the first and third locations (1–30, and 60–90) the probability of arrival of α patients is less than half. In the second and fourth location this probability is greater than half. We use probabilities of 0.7 and 0.3 as before. Once again, starting from a random starting assignment of procedures we track the evolution over 5000 decision dates and plot the frequency of use of A at each location. Procedure A is used almost exclusively in the second and fourth locations, whereas B is used in the first and third. Figure 5 show the details. The emergence of uniform regions is remarkably quick. After that there are small movements at the boundaries, but the regional pattern is very robust.

Finally in Figure 6, we present the outcome from one of our robustness exercises. We make the regions very small — there are forty locations in all, split into four regions. We add noise: after a decision is made by a physician, there is a one in ten chance that she decides to go with the other procedure. We may interpret this as a random component of choice, or a conscious decision to experiment. Despite this large amount of noise, we see that definite regional patterns of procedure use do emerge.

3 Empirical Analysis

3.1 Description of the Data

The data for this study are taken from a census of inpatient stays reported quarterly by Florida hospitals. We selected all patient records for the first half of 1998, subject to several criteria. Patients selected were over 25 years old, had a principle diagnosis of AMI or coronary atherosclerosis, and were admitted either via the emergency room or under a non life-threatening, elective condition.[1] To control for insurance status, we separately analyze patients who are insured by HMOs and Medicare (non-HMO). Finally, for reasons explained below the physician assigned to the patient must have been designated as the attending physician for at least 10 patients at a Florida hospital in the six month period. Resulting samples include 5382 HMO and 5656 Medicare patients and are broadly representative of patients who are candidates for advanced coronary care. A limitation of the data should be noted, however. Each observation is a single hospital stay. Repeat hospitalizations by a single patient are masked in the file, so the complete record of treatment choice is censored. Our econometric model attempts to correct for the possible biases from this censoring.

To test for the significance of social interactions in medical choice, we define the social network for a given physician from the group of physicians who practice in the same hospital(s). That is, we identify the network as the set of other attending physicians who have common admitting privileges in at least one hospital. We exclude from the sample, any physicians treating few (less than 10) patients in order to focus on doctors who are most likely to be interacting in the course of their practice. These networks overlap for individual doctors, rather than forming a partition of the set of doctors. We can view each network as belonging to a given physician, i.e. the *network owner*. For each network, we measure the proportion of all patients treated by network physicians (not including the *owner*) who received each given procedure, including coronary angiography, bypass, and angioplasty. To capture the effect of possible increasing returns to scale, we also include, for each procedure, covariates measuring the total number performed at the hospital. Finally, we use lagged variables to mitigate potential endogeneity between the treatment choice of a physician and that of the network of peers.

[1] The CCS Diagnosis Categories were used to identify the 56 ICD-9CM categories relevant to these patients.

That is, variables measuring the activity by the network physicians are taken from the six months one-year prior to the first half of 1998.

We admit that this construction of social networks is somewhat arbitrary. We considered several alternative variables and methods to define the networks, such as membership in particular medical associations, medical school attended, hospital of residency, and others. Arguably, the hospitals of current activity would be the settings most likely to serve as loci for the types of knowledge spillovers implied in the model. As a check against the specification error in constructing our social network, we consider the effect of randomly-assigned groupings of doctors, to determine whether our network methods serve as a better predictive factor in a given physician's choice. In the next version of the paper we will include the bootstrapped results of empirical models performed using procedure choice rates in randomly constructed physician networks.

Despite the multiplicity of influences that may contribute to any medical decision, we locate procedure choice simply at the level of the attending physician. In the data, the *attending physician*is defined as the party having primary responsibility for a patient's care and treatment, and/or the one who certifies medical necessity. The surgeons who perform bypass or angioplasty operations tend to specialize in just one procedure. Thus, any meaningful choice between the two appears to take place prior to the employment of the surgeon. The attending physician in AMI cases will typically be a cardiologist, to whom a patient may have been referred by a primary care provider.

3.2 Analysis of Angiography

We begin by investigating the binary choice of whether or not to perform coronary angiography[2], a diagnostic procedure that is used to identify, locate and measure the severity of coronary artery disease. While extremely accurate, the procedure is invasive and risky, and may be dominated in terms of cost-effectiveness by other, noninvasive diagnostics such as echocardiography and SPECT (Garber and Solomon, 1999). A probit estimation indicates the factors that affect the decision concerning this procedure, including, possibly, the rate of angiography within the relevant social network.

[2]This procedure is also frequently referred to as cardiac catheterization, although the two are not technically the same thing. Angiography consists of viewing the coronary arteries with an X-ray technique called fluoroscopy. Catheterization is the means by which the dyes required for angiography are delivered to the aorta.

We report probit estimations under each of three models, and for each of two populations distinguished by the insurance type–a group covered by Medicare and a group covered by HMOs. In the first model the only regressors are the district dummies to capture regional variations within 11 planning districts in Florida. In the second model we add patient race, sex, age, and the Charlson index, an age-adjusted index of illness severity summarizing the patient comorbidity characteristics (Charlson, et.al.1987). In the third model, we add the (one-year lagged) social network variables: the proportion of patients treated by others in the physician's network who received the procedure, and the counts of the number of angiographies (*cath count*), bypasses (*bypass count*), and angioplasties (*plasty count*) performed at the hospital in the first half of 1997. Table 1 shows the coefficient estimates from the full model, for each of the populations. Estimation is conducted on separate samples for HMO and Medicare patients to avoid the confounding effects of insurance status on the treatment choices.

Results in Table 1 indicate that the model is broadly consistent with the hypothesis that variations in regions, patient characteristics and social networks are important in the choice of this diagnostic treatment. For example, patients who are admitted from the emergency room or who are transferred to another hospital are less likely to receive the procedure. Older, sicker and black patients are less likely to receive the procedure. Most important for our purpose is the significant effect observed of the average rate of bypass and angioplasty procedures within the physician social network. Other considerations held constant, the transmission of experience via informal contacts with other angiography-prone peers may help explain the readiness to order this diagnostic test. Finally, the coefficients associated with the level of angioplasty and angiography procedures (*plasty count* and *cath count*) at the hospital are consistent with the hypothesis that the economies of scale in performing these procedures may influence the decision, at the margin, to include the angiography procedure.

Table 1 also reports the district dummy coefficients. Further results (not shown) comparing their estimates across models indicate an attenuation of the regional variations in models controlling for patient characteristics and social networks relative to one with regional dummies alone. Perhaps there is a significant role for social networks in explaining some of the variation that initially shows up as regional variation.

Table 1: Probit Model on Coronary Angiography, all variables included.

	HMO patients			Medicare patients		
Variable	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value
Intercept	-0.580	0.163	.0004	1.101	0.275	.0001
emergency	-0.064	0.043	.1359	-0.216	0.044	.0001
transferred	-1.998	0.224	.0001	-2.001	0.336	.0001
district 1	-0.498	0.182	.0064	-0.414	0.131	.0016
district 2	-0.372	0.201	.0636	-0.602	0.147	.0001
district 3	-0.540	0.105	.0001	-0.637	0.114	.0001
district 4	-0.096	0.099	.3330	-0.322	-0.123	.0086
district 5	-0.648	0.096	.0001	-0.580	0.112	.0001
district 6	0.022	0.087	.8004	-0.091	0.101	.3638
district 7	-0.124	0.100	.2121	-0.326	0.114	.0043
district 8	-0.112	0.117	.3369	-0.242	0.114	.0337
district 9	-0.177	0.093	.0557	-0.290	0.108	.0073
district 10	0.055	0.091	.5480	-0.013	0.111	.9058
male	0.077	0.042	.0679	0.063	0.041	.1278
black	-0.390	0.086	.0001	-0.351	0.115	.0022
hispanic	0.100	0.085	.2398	-0.149	0.110	.1755
other race	0.188	0.113	.0964	0.017	0.057	.9158
age	-0.011	0.002	.0001	-0.031	0.003	.0001
charlson	-0.071	0.020	.0003	-0.123	0.020	.0001
bypass rate	1.914	0.159	.0001	1.714	0.155	.0001
plasty rate	1.803	0.157	.0001	2.486	0.164	.0001
bypass count	0.011	0.045	.8003	0.075	0.044	.0858
plasty count	-0.753	0.093	.0001	-0.966	0.095	.0001
cath count	1.211	0.099	.0001	1.389	0.099	.0001

N(HMO)= 5382; N(Medicare)= 5656, Log likelihoods: $(-2856.36, -2728.83)$

3.3 Two-stage Treatment Model

Patients who are hospitalized with AMI and coronary atherosclerosis may be given intensive surgical interventions such as bypass surgery or angioplasty, or may simply be held for observation, diagnostic testing and drug therapy. The data reveal that the two surgical interventions plus non-surgery inpatient stays account for the preponderance of patient care given in our data. It seems reasonable to posit a two-stage process in patient care decisions. First, a decision is made about whether or not to intervene surgically. Perhaps at this stage, a decision to delay surgery for a future hospitalization is made or the surgery is not recommended for the patient. For other patients, one of the two principle surgical methods is prescribed. Thus, conditional on the first decision, a choice between bypass and angioplasty is necessary for those selected for surgery. Whether treatment decisions are influenced by the norms of practitioners in the physician's social network may be modeled with two related equations. The events of interest are the determination of surgical intervention and the choice of bypass surgery versus angioplasty, given that surgery will be given the patient.

We use a bivariate probit model of treatment choices. In the data, we observe the qualitative outcome $y_1 = 1$ if the patient receives surgery, and $y_1 = 0$ otherwise. The type of surgery is observed as $y_2 = 1$ if bypass and $y_2 = 0$ otherwise. Latent (unobserved) variables determine the patient's illness and location characteristics or preferences for surgery, y_1^*, and the type of surgery, y_2^*, while only the qualitative choices of (y_1, y_2) are observed. The two equations are:

$$y_1^* = X\beta_1 + \epsilon_1$$

$$y_2^* = X_2\beta_2 + \epsilon_2$$

where ϵ_1, ϵ_2 are assumed to be i.i.d. bivariate normal random variables with zero means and finite variances and covariance. The choice of surgery is a function of exogenous variables including additional ones required for identification, $X = (X_1, X_2)$. A consistent estimator of the parameters of the model requires a further assumption about the data. The qualitative outcome $y_1 = 1$ is observed when $y_1^* > 0$, and $y_1 = 0$ may correspond to $y_1^* \leq 0$. However, it is also possible that when $y_1 = 0$, $y_1^* > 0$, but the surgery is simply delayed and not observed. Thus, we estimate a Heckman probit model with selection in which y_2 is only observed when outcome $y_1 = 1$. Thus, we explain the choice between bypass and angioplasty with selection determined on whether the patient received surgery or not.

The results are reported in Table 2. The HMO and Medicare samples are again separated in estimation to avoid the confounding effects of insurance status on the treatment choices. In the model, the selection equation for y_1 determines who is considered for the bypass or angioplasty. This first stage decision is partly censored since patients discharged without surgery may be scheduled for later. The model shows that the decision to withhold expensive surgical procedure is correlated with patients admitted from emergency and those who are discharged as a transfer to another hospital. Patient characteristics capturing comorbidity, gender, age and race are also significant. The tendency to promote surgery in hospitals that do relatively large numbers of surgeries is partly confirmed by the significant coefficients on the counts of bypass and coronary angiographies, but not for the volume of angioplasty. But after controlling for these factors and the regional variations, there remain in the evidence substantial effects associated with the propensity to choose surgery within the physician's social network. The rates of angioplasty in the network lead to higher rates of surgery and of bypass relative to angioplasty. The positive inducement to surgery brought on by physician interaction is confirmed separately in both samples of patients.

4 Conclusion

Our study of coronary patients, like many others, finds variations in the rates of diagnostic and treatment choices for patients with heart attacks or coronary atherosclerosis. These variations are sustained, even between regions of a single state, after controlling for demographic and illness conditions in some detail.

Observing these variations in the utilization of expensive methods of treatment or diagnosis of coronary diseases does not necessarily imply unwanted or welfare-reducing choices. Based on the results of our theoretical model, we argue that the patterns and relationships we observe in the data may reflect an equilibrium process that manifests the influence of a local interaction structure. As we have stated above, the influence of such interactions may derive from knowledge spillovers across physicians, or from a taste for conformity, or both. These effects are distinct from economies of scale effects that reward high volume production of these procedures.

To examine social interaction among physicians in hospital settings for advanced treatment such as this one, we traced the aggregate treatment

Table 2: Bivariate probit with sample selection: two-stage surgery choice.

	HMO patients			Medicare patients		
Variable	Estimate	Std. Error	P-value	Estimate	Std. Error	P-value
Bypass vs. Angioplasty (y_2)						
emergency	-0.413	0.067	.0001	-0.143	0.077	.0611
district 1	0.058	0.207	.7781	-0.238	0.170	.1619
district 2	0.432	0.229	.0592	1.068	0.180	.0001
district 3	0.591	0.128	.0001	1.134	0.150	.0002
district 4	0.591	0.128	.0023	0.419	0.156	.0068
district 5	0.260	0.120	.0300	0.609	0.145	.0001
district 6	0.102	0.105	.3312	0.225	0.132	.0883
district 7	0.020	0.117	.8618	0.166	0.147	.2614
district 8	0.553	0.137	.0003	0.863	0.152	.0001
district 9	0.314	0.116	.0070	-0.034	0.150	.8189
district 10	0.215	0.113	.0571	0.240	0.148	.1041
male	0.232	0.050	.0001	0.222	0.049	.0001
black	-0.254	0.116	.0293	0.013	0.153	.9310
hispanic	-0.059	0.101	.5569	0.002	0.146	.9877
other race	0.037	0.120	.7556	0.277	0.181	.1263
age	-0.005	0.003	.1111	-0.005	0.005	.2862
charlson	0.206	0.024	.0001	0.072	0.024	.0028
bypass rate	-0.540	0.214	.0123	-1.891	0.256	.0004
plasty rate	2.494	0.243	.0003	2.534	0.314	.0001
bypass count	0.346	0.049	.0001	0.405	0.052	.0001
plasty count	-0.934	0.113	.0001	-0.752	0.121	.0001
cath count	0.590	0.124	.0002	0.285	0.130	.0279
intercept	-1.799	0.225	.0001	-1.103	0.381	.0040
Surgery or not (y_1)						
emergency	-1.540	0.070	.0001	-1.671	0.069	.0001
transferred	-2.244	0.266	.0001	-2.619	0.521	.0001
district 1	-0.164	0.256	.5227	-0.111	0.177	.5289
district 2	0.032	0.275	.9092	0.121	0.205	.5560
district 3	0.059	0.141	.6776	-0.194	0.154	.2071
district 4	0.409	0.134	.0020	0.172	0.160	.2832
district 5	-0.620	0.130	.0001	-0.561	0.152	.0001
district 6	0.095	0.117	.4189	-0.118	0.131	.3658
district 7	-0.211	0.143	.1412	-0.128	0.157	.4153
district 8	-0.120	0.146	.4122	-0.061	0.141	.6669
district 9	0.105	0.120	.3799	-0.050	0.139	.7191
district 10	-0.020	0.124	.8728	-0.189	0.149	.2030
male	0.318	0.057	.0001	0.203	0.057	.0001
black	-0.541	0.102	.0001	-0.590	0.144	.0004
hispanic	-0.026	0.113	.8173	-0.256	0.137	.0617
other race	0.162	0.183	.3749	-0.021	0.219	.9230
age	-0.025	0.003	.0001	-0.062	0.004	.0001
charlson	0.001	0.026	.9560	-0.131	0.028	.0001
bypass rate	1.453	0.203	.0001	1.807	0.195	.0001
plasty rate	2.272	0.206	.0002	3.025	0.219	.0003
bypass count	0.438	0.091	.0004	0.556	0.085	.0001
plasty count	-0.181	0.144	.2100	-0.193	0.140	.1679
cath count	0.641	0.146	.0001	0.415	0.152	.0060
intercept	1.168	0.216	.0001	4.407	0.373	.0001
rho	0.374	0.131		0.046	0.120	

N(HMO)=5382, N(Medicare)=5656; Log-likelihood(HMO)= -3546.079, Wald(HMO)=419.50; Log-likelihood(Med.)=-3278.813,Wald(Med.)=394.12

tendencies of a hypothetical social network by identifying the most likely points of mutual contact between physicians. In doing so, we have found that, controlling for patient characteristics, a patient will be more likely to receive angiography or the surgical options if the attending physician is in a group prone to recommend those options. While our construction of social networks is merely suggestive of what in reality is a much richer and more subtle set of interactions, it contains information that is empirically relevant to treatment choices, and its explanatory power appears robust across patient populations.

Appendix

Proof of Proposition 1

Proof. Since the process on Z is irreducible and aperiodic, it has an essentially unique invariant distribution. Each state can be specified in terms of m, the location of the first zero. First we define the probabilities $b(m)$ and $d(m)$ of transition $m \to m+1$ and $m \to m-1$ respectively. Recalling that the rate of arrival of patients is one, these are given by:

$$b(m) = \begin{cases} p_W & \text{if } m < 0 \\ \\ p_E & \text{otherwise.} \end{cases}$$

In other words, m moves to the right if an α-patient arrives at m, which happens with probability p_W in the West and p_E in the East.

$$d(m) = \begin{cases} 1 - p_W & \text{if } m \leq 0 \\ \\ 1 - p_E & \text{otherwise.} \end{cases}$$

In other words, m moves to the left if a β-patient arrives at $m-1$, which happens with probability $1 - p_W$ in the West and $1 - p_E$ in the East.

The process is reversible, so that invariant distributions can be obtained from the detailed balance conditions:

$$b(m-1)\rho(m-1) = d(m)\rho(m).$$

We can confirm that these are satisfied. In case $m \leq 0$, we can substitute for ρ and confirm that

$$\frac{b(m-1)}{d(m)} = \frac{p_W}{1-p_W} = \frac{\rho(m)}{\rho(m-1)}.$$

When $m > 0$,

$$\frac{b(m-1)}{d(m)} = \frac{p_E}{1-p_E} = \frac{\rho(m)}{\rho(m-1)}.$$

So $\rho(\cdot)$ is an invariant distribution. It is not a mixture of δ_0 and δ_1, hence the process coexists. $\square$

References

1. Charlson, M.E., Pompei P., Ales K.L., McKenzie C.R. (1987), A new method of classifying prognostic comorbidity in longitudinal studies: development and validation, *Journal of Chronic Disorders,* 40(5), 373–383

2. Ellison, G. (1993), Learning, Local Interaction, and Coordination, *Econometrica,* 61, 1047–1071.

3. Ellison, G. and D. Fudenberg (1993), Rules of Thumb for Social Learning, *Journal of Political Economy,* 101, 612–643.

4. Elixhauser *et al.* (1998), Comorbidity Measures for Use with Administrative Data, *Medical Care,* 36, 8–27.

5. Fournier, Gary M. and Melayne Morgan McInnes, (2002), The Effects of Managed Care on Medical Referrals and the Quality of Specialty Care, *Journal of Industrial Economics.*

6. Garber, Alan M. and Neil A. Solomon (1999), Cost-Effectiveness of Alternative Test Strategies for the Diagnosis of Coronary Artery Disease, *Annals of Internal Medicine,* 130, 719-728

7. Health Care Financing Administration, *Clinical Classifications Software (CCS), 2002 Software and User's Guide.*

8. Kandori, M., G. Mailath and R. Rob (1993), Learning, Mutation, and Long Run Equilibria in Games, *Econometrica,* 61, 29–56.

9. Morris, S. (2000), Contagion, *Review of Economic Studies,* 67(1), 57-78

10. Phelps, C. E. (2000), Information Diffusion and Best Practice Adoption, in: Handbook of Health Economics, Volume 1, Edited by A.J. Culyer and J.P. Newhouse, Elsevier Science B.V., pp.223–264.

11. Romano P.S., Roos L.L., Jollis J. (1993), Adapting a clinical comorbidity index for use with ICD-9-CM administrative data: Differing perspectives, *Journal of Clinical Epidemiology,* 46(10), 1075–1079.

12. Young H.P., and M.A. Burke (2001), Competition and Custom in Economic Contracts: A Case Study of Illinois Agriculture, *The American Economic Review,* 91, 559–573.

13. Young, H. P. (1993), The Evolution of Conventions, *Econometrica,* 61, 57-84.

14. Young, H. P. (1996), The Economics of Convention, *Journal of Economic Perspectives*, 10, 105-122.

Figure 1. Frequency of use of procedure *A* at location *m*

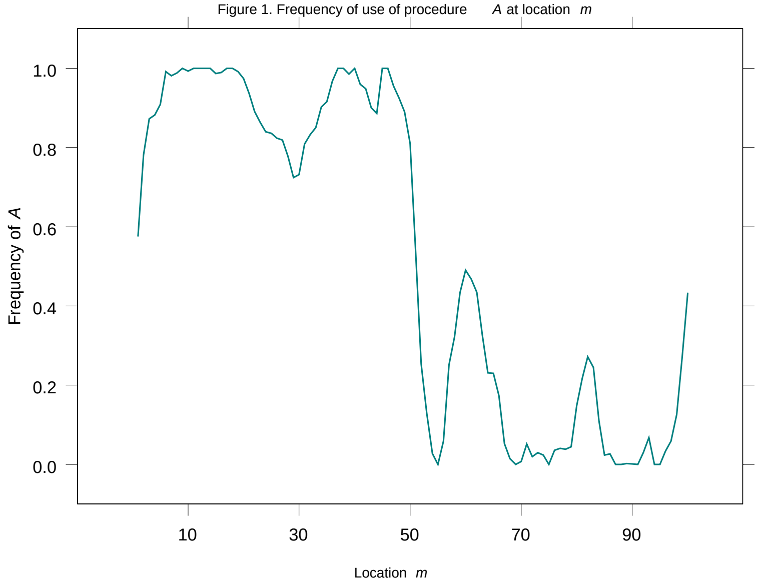

Figure 2. Use of *A* over time

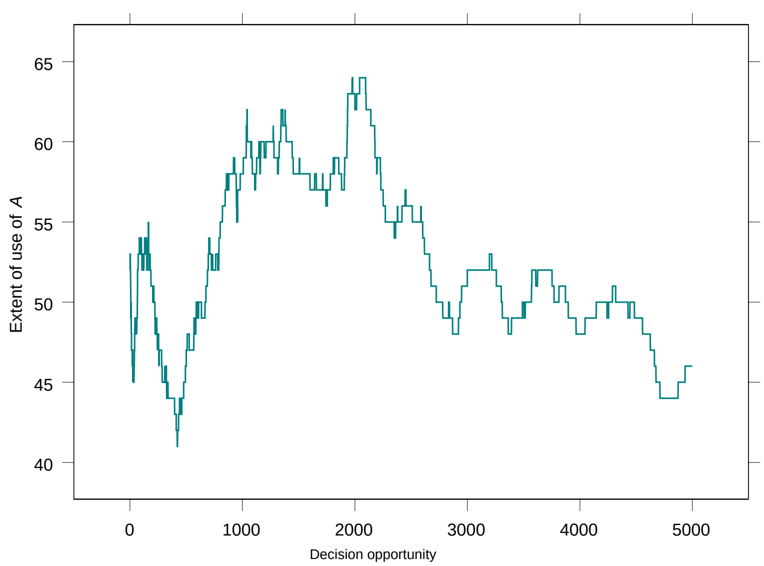

Figure 3. Choice of Procedure at Locations

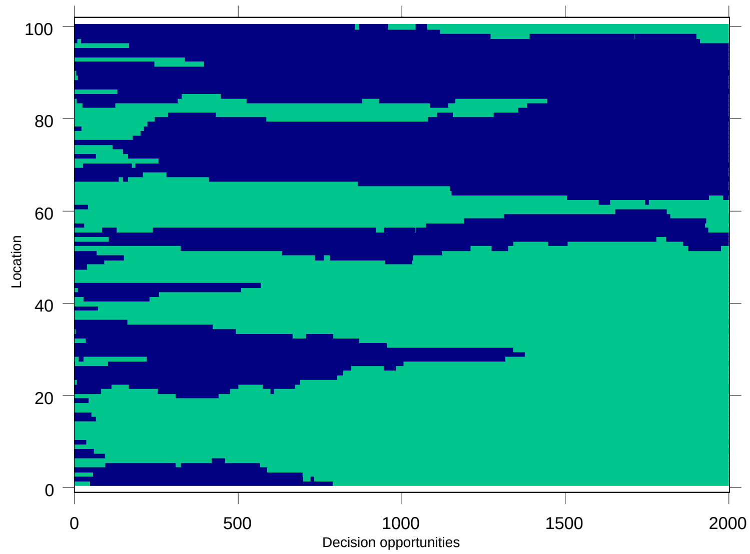

Figure 4. Frequency with which procedure *A* is used at location *m*

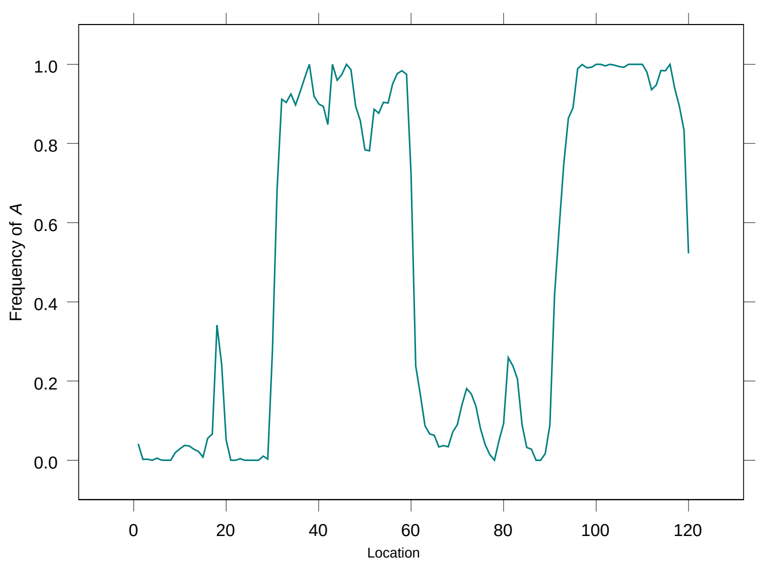

Figure 5. Choice of Procedure at Locations

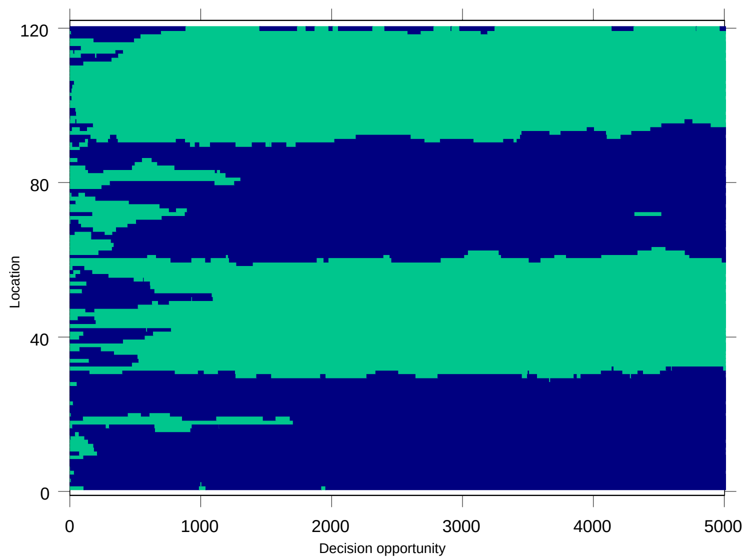

Figure 6. Small groups, large noise